

IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal LetterASSISTANT COMMISSIONER FOR PATENTS
Washington, D.C. 20231

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility ☐ Design☒ original patent application,
☐ continuation-in-part application

INVENTOR(S): Qinghua ZHENG

TITLE: METHOD AND APPARATUS FOR VALIDATING CROSS-ARCHITECTURE ISA EMULATION

Enclosed are:

- ☒ The Declaration and Power of Attorney. ☒ signed ☐ unsigned or partially signed
☒ 2 sheets of drawings (one set) ☐ Associate Power of Attorney
☐ Form PTO-1449 ☐ Information Disclosure Statement and Form PTO-1449
☐ Priority document(s) ☐ (Other) (fee \$)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	20 — 20	0	\$18	\$ 0
INDEPENDENT CLAIMS	3 — 3	0	\$78	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$260	\$ 0
BASIC FEE: Design \$310.00 ; Utility \$690.00				\$ 690
TOTAL FILING FEE				\$ 690
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 690

Charge \$ 690 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. _____

Date of Deposit _____

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231.

By _____

Typed Name:

Respectfully submitted,

Qinghua ZHENG

By

John K. Harrop

Attorney/Agent for Applicant(s)

Reg. No. 41,817

Date: April 13, 2000

Telephone No.: (202) 824-8821

METHOD AND APPARATUS FOR VALIDATING CROSS-ARCHITECTURE ISA EMULATION

Technical Field

The technical field is emulation of computer instruction sets.

Background

Cross-architecture emulation is needed when running native applications on a target platform, or computer, that uses an instruction set architecture (ISA) different from the ISA for which the native applications were initially intended. The developer of a cross-architecture emulation product wants to know that the emulator correctly translates or interprets the native applications. Various schemes exist to verify cross-architecture emulation. However, these schemes cannot easily, if at all, comprehensively test the emulation process.

To improve the emulation process, a binary translation product that runs on the target platform may be used to emulate native instructions running on a native, or legacy, platform. Binary translation automatically translates binary code from the native instruction set to binary code for the target platform without the need for high-level source code. As with any emulation product, a developer of the binary translation product may desire to verify the accuracy of emulation from the native ISA to an ISA operating on a target platform.

One conventional approach to verifying cross-architecture emulation is to send as many native applications as possible through the binary translation product, or emulator, and then verify that the outputs produced by the binary translation product are identical to corresponding outputs produced by the native applications running on the native platform. This conventional approach has several disadvantages. First, the binary translation product designer cannot know whether testing is complete because the native applications are usually compiler-generated and the machine code generated by the binary translation product only includes a subset of the instruction set architecture. Machine instructions that are not in this subset will never be generated in the compiler so that some binary instructions are never tested through the emulation process. Second, running an application through the binary translation product on the target platform and running the same application through the native platform is cumbersome. A separate process may then be needed to verify the results. Third, when an error is detected, pinpointing the

exact machine instruction that caused the error may be difficult or impossible. In addition, replication of the emulation error may be impossible to achieve. Fourth, execution of some applications is time- or system-sensitive such that the result of an execution may not be reproducible, which adds to the difficulty of verifying binary emulation.

Summary

A method and an apparatus allows complete and efficient verification of cross-architecture ISA emulation. A random verification framework runs concurrently on two different computer architectures. The framework operates without regard to existing native applications and relies instead on binary instructions in a native ISA. The framework is able to determine emulation errors at a machine instruction level. The emulation errors are then easily identifiable and reproducible.

A random code generator generates one or more sequences of native machine instructions and corresponding initial machine states in a pseudo-random fashion. The native instructions are generated from an entire set of the native ISA. The instructions and the state information are provided to initialize a native computer architecture. The same instructions and state information are provided using standard machine-to-machine languages, such as TCP/IP, for example, to a target computer architecture. The target computer architecture may be embodied as an actual hardware device, or may be a simulation of a yet-to-be-built computer architecture. The target computer architecture includes a binary emulator that translates the native instructions into binary instructions executable on the target computer architecture. The final states of the native and the target computer architectures are gathered, and a verification engine compares the results. Any differences may indicate an emulation error or failure.

The random verification framework may be run continuously to test emulation of the complete set of instructions from the native ISA. Even the least-used machine instructions are tested by the framework. Further, the framework automates the emulation verification process. Inter-machine communications allow the native and the target computer architectures to process the same machine instructions from the same initial states. Any inconsistencies in the final produced states indicate the emulation

error. The framework can then easily pinpoint the exact machine instruction, register number and input machine state that caused the emulation error, thereby significantly reducing the amount of time required for debugging. Finally, the emulation errors detected using this framework are easily reproducible.

Description of the Drawings

The detailed description will refer to the following drawings, in which like numerals refer to like objects, and in which:

Figure 1 is a block diagram of an example of an apparatus for validating cross-architecture ISA emulation; and

Figure 2 is a flowchart illustrating an operation of the apparatus of Figure 1.

Detailed Description

A method and an apparatus allow complete and efficient verification of cross-architecture ISA emulation. A random verification framework runs concurrently on two different computer architectures. The framework is able to determine emulation errors at a machine instruction level. The emulation errors are then easily identifiable and reproducible.

Figure 1 is a block diagram of an embodiment of an apparatus for validating cross-architecture instruction set architecture (ISA) emulation. A framework 10 includes a random code generator (RCG) 20 that generates native machine instruction level code and an initial machine state. The machine instruction level code (i.e., sequences of binary instructions) and initial machine state are provided to a native ISA platform 11, which includes a first, or X, execution engine 30. The X execution engine 30 is capable of executing the instruction level code without emulation or translation.

The RCG 20 may include a probability file that is used to pseudo-randomly generate the machine instruction level code. The RCG 20 also provides the machine instruction level code and the initial machine state to a non-native ISA or target platform 12, which includes a second, or Y, execution engine 40. To execute the machine instruction level code, the Y execution engine 40 emulates or translates the machine instruction level code using an emulator 45.

1 In an alternative embodiment, the emulator 45 may operate on a target platform
2 simulator (not shown), which in turn operates on the native platform 11.

3 The X execution engine 30 and the Y execution engine 40 may execute the same
4 machine instruction level code concurrently. The execution engines may perform
5 information transfer using any standard machine-to-machine protocol, such as TCP/IP,
6 for example.

7 The X execution engine 30 and the Y execution engine 40 each will provide a
8 final machine state. The final machine states are sent to a verification engine 50. The
9 verification engine 50 compares the final machine states to determine any differences.
10 Such differences may indicate an error in emulation of the native machine instruction
11 level code. In particular, the verification engine 50 is able to pinpoint the exact machine
12 instruction, register number and input machine state that caused the emulation error. The
13 exact source of the emulation error may be particularly easy to locate because the
14 instruction sequence that generated such an emulation error is typically a short sequence
15 of binary instructions. Hence, a review of information saved when an emulation failure
16 occurs allows the test designer to quickly pinpoint the exact cause of the emulation
17 failure.

18 In an example, if a binary instruction to be emulated comprised ADD R1 and R2,
19 R3, the verification engine 50 would determine if the result written to the register R3
20 using the emulator 45 and the Y execution engine 40 was the same as that written to the
21 register R3 using the X execution engine 30. As long as the final state produced by the
22 X execution engine 30 is the same as the final state produced by the Y execution engine
23 40, the emulation of the binary instruction with a given specific initial machine state was
24 correct.

25 The framework 10 relies on pseudo-random generation of machine level
26 instructions to comprehensively test the correct emulation of native applications on a
27 second computer architecture. The machine level instructions may be generated
28 according to many different random generation schemes. In an embodiment, each
29 machine level instruction is assigned a probability. In particular, a specific instruction
30 to be generated is controlled by an input probability file, which is a list of pre-defined

machine instructions, each with an associated probability. The machine instructions are arranged in a hierarchy form and divided into segments based on a function of the instruction. A first such segmentation may include floating instructions and CPU instructions. Continuing, the CPU instructions may be further segmented according to arithmetic/logic, immediate, memory, system, memory management, and branch instructions, and other CPU instructions. The arithmetic/logic instructions may be segmented into ADD, SUB, AND, OR, and XOR instructions, and other arithmetic/logic instructions, for example. Each hierarchical instruction is assigned a probability such that a cumulative probability along any hierarchical path equals 1.0.

The instruction sequence generation process is pseudo-random because any random code generator must operate according to a pre-determined algorithm. Further limitations also lead to a loss of true randomness. For example, certain memory regions may be inaccessible to the random code generator.

An example probability file that controls the pseudo random instruction generation follows. In the example, expressions such as `pr_xxx` refer to a probability value for binary instruction `xxx`.

The instructions are arranged in hierarchical segments.

```
pr_cpu                                = 1.00
pr_fp                                = 0.00

pr_cpu                                = 1.00
pr_cpu_arithlog                      = 0.25
pr_cpu_immediate                     = 0.25
pr_cpu_shexdet                       = 0.00
pr_cpu_mem                           = 0.50

.
.
.

pr_cpu_arithlog                      = 0.25
```

1	pr_add	= 0.04
2	pr_add1	= 0.04
3	.	
4	.	
5	.	
6	pr_sub	= 0.04
7	pr_subb	= 0.04
8		
9	pr_cpu_immediate	= 0.25
10	pr_ldo	= 0.15
11	.	
12	.	
13	.	
14	pr_subi	= 0.10
15	.	
16	.	
17	.	

With the above-assigned probabilities, the RCG 20 will never select a floating point (fp) instruction (i.e., pr_fp = 0.00). Note that the cumulative probability of the first segment is 1.00. That is, pr_cpu and pr_fp sum to 1.00. This cumulative probability rule is maintained throughout the probability file.

By assigning higher probability values to some instructions and lower probability values to other instructions, the test designer can ensure that all binary instructions are eventually tested.

The RCG 20 begins with a seed value and then determines a pseudo random probability value. For example, a seed value of 10 may produce a probability sequence of 0.34, 0.8, Using the cumulative probability in the above example, a random probability of 0.34 would lead the RCG 20 to generate a cpu_immediate instruction, and a random probability of 0.8 would lead the RCG 20 to generate a pr_cpu_immediate_subi instruction.

Figure 2 is a flowchart of a process that may be used by the framework 10 of Figure 1. The process includes a main process and sub-processes A and B. Sub-process A operates on the native platform 11 and the X execution engine 30. Sub-process B may operate in one of at least two ways. First, sub-process B may operate on the binary emulator 45, which may in turn operate on the target platform 12. Second, sub-process B may operate on the binary emulator 45, but the binary emulator 45 may operate on a target platform simulator. The target platform simulator may then operate on the native platform 11.

The main process begins at 100. In code generate block 110, the RCG 20 generates pseudo-random native machine code and generates a random initial machine state. The native machine code and the initial machine state are provided to the native platform 11, and the native platform 11 is initialized using the initial machine state, block 120. The same native machine code and initial machine state are also provided to the target platform 12, and the target platform 12 is initialized, block 130. The native machine code and initial machine state may be provided to the target platform 12 using TCP/IP protocols, for example.

In block 140, the X execution engine 30 executes the native machine code. Concurrently, the binary emulator 45 emulates the entire Y execution engine 40, which internally executes the native machine code, block 150. Next, a final machine state S1 is collected in the native platform 11, block 160, and a final machine state S2 is collected in the target platform 12, block 170.

In block 180, the target platform 12 provides the final machine state S2 to the verification engine 50, using TCP/IP protocols, for example. The verification engine 50 compares the final machine states S1 and S2. If the final machine states S1 and S2 are equal, the process moves to block 220 and ends. If the final machine states S1 and S2 are not equal, the process moves to block 210, and information related to the emulation failure (i.e., the initial and final states, and the instruction sequence) are written to a file. The process then moves to block 220 and ends.

The test designer then need only refer to the file to determine the source of the emulation failure, and to reproduce such emulation failure.

1 The terms and descriptions used herein are set forth by way of illustration only
2 and are not meant as limitations. Those skilled in the art will recognize that many
3 variations are possible within the spirit and scope of the invention as defined in the
4 following claims, and there equivalents, in which all terms are to be understood in their
5 broadest possible sense unless otherwise indicated.

In The Claims:

1. An apparatus for pseudo-random testing binary emulation, comprising:
 - a random code generator that generates an initial machine state and a binary instruction sequence;
 - a native architecture execution engine that executes the binary instruction sequence to produce a final state S1;
 - a target architecture execution engine that executes the binary instruction sequence to produce a final state S2, the target architecture execution engine comprising a binary emulator that emulates the binary instruction sequence according to the target architecture; and
 - a verification engine that compares the final state S1 and the final state S2, wherein when the final state S1 and the final state S2 do not match, an emulation failure has occurred.
2. The apparatus of claim 1, wherein the native and the target architecture execution engines communicate using machine-to-machine communications protocols.
3. The apparatus of claim 2, wherein the communications protocol is a TCP/IP protocol.
4. The apparatus of claim 1, wherein when the emulation failure occurs, information related to the emulation failure is written to a file.
5. The apparatus of claim 1, wherein the target platform is a simulator.
6. The apparatus of claim 1, wherein the target platform is a hardware embodiment.
7. The apparatus of claim 1, wherein the random code generator comprises a probability file comprising probability values for each instruction in a native instructions set architecture (ISA).
8. The apparatus of claim 7, wherein the probability values in the probability file are user-generated.
9. A method for pseudo-random testing of binary emulation, comprising:
 - generating a pseudo-random binary instruction sequence;
 - generating an initial machine state;
 - initializing a native machine according to the initial machine state;

- 1 executing the binary instruction sequence on the native machine to produce a final
2 state S1;
3 initializing a target machine according to the initial machine state;
4 emulating the binary instruction sequence on the target machine;
5 executing the emulated binary instruction sequence to produce a final state S2;
6 and
7 comparing the final state S1 to the final state S2 to determine an emulation error.
- 8 10. The method of claim 9, further comprising:
9 communicating the initial machine state and the binary instruction sequence to
10 the target machine using a machine-to-machine communication protocol; and
11 communicating the final state S2 to the native machine using the machine-to-
12 machine communication protocol.
- 13 11. The method of claim 10, wherein the machine-to-machine communication
14 protocol is a TCP/IP protocol.
- 15 12. The method of claim 9, further comprising storing information related to the
16 emulation failure, the information comprising the initial machine state, the binary
17 instruction sequence and the final states S1 and S2.
- 18 13. The method of claim 9, wherein the binary emulation is executed on a simulator.
- 19 14. The method of claim 9, wherein the binary emulation is executed on a hardware
20 device.
- 21 15. The method of claim 9, further comprising:
22 assigning a probability value to each binary instruction in a native instruction set
23 architecture (ISA);
24 pseudo-randomly generating a probability value; and
25 selecting the binary instruction sequence based on the probability value.
- 26 16. A method for verifying cross-architecture emulation, comprising:
27 randomly generating a native instruction sequence and an initial machine state;
28 providing the native instruction sequence and the initial machine state to a first
29 platform having a first instruction set architecture (ISA);

1 providing the native instruction sequence to a second platform having a second
2 ISA;
3 emulating the native instruction sequence according to the second ISA;
4 executing the native instruction sequence in the first platform, the execution
5 providing a final state S1;
6 executing the emulated native instruction sequence on the second platform, the
7 execution providing a final state S2; and
8 comparing the final states S1 and S2 to determine an emulation failure.

9 17. The method of claim 16, wherein the native instruction sequence is randomly
10 generated from a set of all instructions in the first ISA.


11 18. The method of claim 16, wherein the random generation is based on a user-
12 defined value assigned to each instruction in the first ISA.

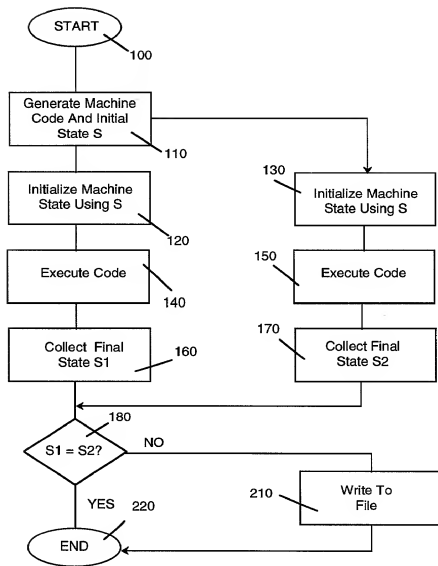
13 19. The method of claim 16, wherein the emulation of the native instruction sequence
14 and the execution of the emulated native instruction sequence is completed on a binary
15 instruction emulator operating on a simulator which operates on the first platform.

16 20. The method of claim 16, wherein the emulation is completed on a binary
17 instruction emulator operating on the second platform.

Abstract

A method and an apparatus allows complete and efficient verification of cross-architecture ISA emulation. A random verification framework runs concurrently on two different computer architectures. The framework operates without regard to existing native applications and relies instead on binary instructions in a native ISA. The framework determines emulation errors at a machine instruction level. A random code generator generates one or more sequences of native machine instructions and corresponding initial machine states in a pseudo-random fashion. The native instructions are generated from an entire set of the native ISA. The instructions and the state information are provided to initialize a native computer architecture. The same instructions and state information are provided using standard machine-to-machine languages, such as TCP/IP, for example, to a target computer architecture. A binary emulator then translates the native instructions so that the instructions may be executed on the target computer. Alternatively, the binary emulator may be embodied as a software routine operating on a simulator, which in turn operates on the native computer architecture. The final states of the native and the target computer architectures are gathered, and a verification engine compares the results. Any differences may indicate an emulation error or failure. The random verification framework may be run continuously to test emulation of the complete set of instructions from the native ISA.

*Fig. 1*

*Fig. 2*

DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION

ATTORNEY DOCKET NO. 10001205-1

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

METHOD AND APPARATUS FOR VALIDATING CROSS-ARCHITECTURE ISA EMULATION

The specification of which is attached hereto unless the following box is checked:

() was filed on _____ as US Application Serial No. or PCT International Application Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES _____ NO _____
			YES _____ NO _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879

Place Customer
Number Bar Code
Label hereSend Correspondence to:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80528-9599**Direct Telephone Calls To:**Alan H. Haggard
(408) 447-0077

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Qinghua ZHENGCitizenship: ChinaResidence: Campbell, CaliforniaPost Office Address: 81 Shelley Avenue, Campbell, California 95008

Inventor's Signature

Date